DOCUMENT RESUME

ED 302 183 IR 013 542

AUTHOR Weyer, Stephen A.
TITLE As We May Learn.
PUB DATE Jun 86
NOTE 21p.; Paper presented at an Invitational Conference on Multimedia in Education (Cupertino, CA, June 19-20, 1986). For the complete journal, see IR 013 538.
PUB TYPE Journal Articles (080) -- Speeches/Conference Papers (150)
JOURNAL CIT Learning Tomorrow: Journal of the Apple Education Advisory Council; n3 p89-109 Spr 1987

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS Computer System Design; *Expert Systems; *Information Retrieval; Information Systems; Learning Strategies; *Man Machine Systems; *Multimedia Instruction; *User Needs (Information)
IDENTIFIERS Information Management; *Intelligent Tutoring Systems; Knowledge Bases

ABSTRACT

Using as a backdrop Vannevar Bush's 1945 discussion of the "Memex," his proposed system for organizing and storing information, this paper considers the development of intelligent multimedia knowledge systems. Particular attention is focused on learner needs and intentions in information seeking, and the implications of these factors for the development of intelligent multimedia knowledge systems are considered. The importance of a user-oriented information system is stressed. The paper concludes with a discussion of the possibility of creating hypermedia systems that go beyond the improvement of searching and browsing techniques to create intelligent systems that might almost emulate expert human behavior. (3 figures and 11 references) (EW)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☒ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy

ED302183

As We May Learn

Stephen A. Weyer
Apple Computer, Inc.

BEST COPY AVAILABLE

IR013542

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY
Saul R[illegible]kman
TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

Vannevar Bush's classic article of 40 years ago, "As We May Think," [Bush45] charted a challenging peacetime course for information researchers. The problem then was an increasingly specialized "growing mountain of research" that, ironically, might hinder further scientific progress because of difficulties in sifting through this mass of material. His proposed "Memex" would have been a library-in-a-desk with typed items, photographs, and longhand annotations recorded on microfilm. More than just an extensive and efficient storage medium, however, it would also have provided powerful ways to consult and modify the human record, accessible by association as well as by indexing. Any two items could be joined together by the user, including the user's own materials and notes. Chains of these associations would form a "trail" with many possible side trails. Trails could be named and shared with other information explorers.

"Wholly new forms of encyclopedias will appear, ready-made with a mesh of associative trails running through them, ready to be dropped into the Memex and there amplified." In addition to scientists, lawyers, patent attorneys, chemists, and historians, "there is a new profession of trailblazers, those who find delight in the task of establishing useful trails through the enormous mass of the common record. The inheritance from the master becomes not only his additions to the world's record, but for his disciples the entire scaffolding by which they were erected."

In closing, Bush ponders not only whether a Memex may be necessary for our complex civilization to survive and understand itself, but also whether the media of storage and mode of interaction might more directly involve the brain, bypassing the tactile, oral, and visual levels of representation. Neural input and output, as replacements for multimedia user interfaces, still seem distant possibilities (except in science fiction novels such as Vernor Vinge's *True Names* or William Gibson's *Neuromancer*).

In addition to its filing-cabinet functions, there were hints in its design of being able to use the Memex for numerical and logical manipulations. With the benefit of hindsight, we might have asked Bush to add further symbolic processing, so that an automatic file clerk of sorts might have assisted us in searching, organizing, and interpreting the data. We might also have asked: if Memex became indispensable to the work of his professional scientists, how might it revolutionize the work of amateur scientists, artists, and learners of all ages and abilities?

With the advent of inexpensive, powerful personal workstations and disk storage, CD-ROM and WORM optical storage, high-bandwidth networks and hypertextlike software systems, much of Bush's vision has begun to materialize. While building and experimenting with such systems, it's important to begin looking beyond Memex and beyond browsing, not only because browsing breaks down when there are too many trails (this point is elaborated later in the paper) but also because the importance of individual trails will often relate to educational purpose; instead of hypertext systems, we need to consider "hyperknowledge assistants." This paper's title (which I hope is not too pretentious in its extension of Bush's title) is meant to draw attention to the educational potential of such a system—"learning" in addition to "thinking." The name "Memex" seems to connote memory and storage; perhaps we should call our system "Mimex" to emphasize simulation and learning from examples.

The basic theme of this paper is that the system, rather than being only a passive tool or repository for information, should become a partner with the learner in searching, interpreting, and creating knowledge. As important, we should look at information from the perspective of the learner and his/her needs, rather than by what the technology can do or what the information industry currently provides. Considering such symbiotic systems and information needs should help us to rethink many of our assumptions about electronic information systems. We need fundamentally new organizing principles for knowledge, and new navigation and manipulation tools for the learner.

Intelligent Multimedia Knowledge Systems for Learning

In this section, I will touch briefly on the importance of the aspects of *multimedia, learning*, and *intelligence* in broadening our perspective on information systems. Most current information systems deal only with text; they retrieve information but provide no help in analysis or synthesis, and treat the user's query and its stored data and information as "just bits" rather than as "meaning." Literal retrieval and exhaustive browsing are sometimes necessary, but more often the learner will want to define similarity in more personal or intuitive ways and to filter and conform the information to suit his or her own needs.

Multimedia sensations are essential in characterizing and understanding the world around us— perception and experience are highly dynamic and personal, not well portrayed by static text. To provide flexible access, a knowledge base must describe textual, visual, and aural images with more than a name—identifying and classifying important characteristics. In addition to requirements for the representation of content, a multimedia perspective should change the style in which we present and organize knowledge—how text, animated pictures, music, and speech are combined to explain a principle of physics or organize a roadmap to a new subject area. To support access for the needs of presentation, the knowledge base must include advice related to when to use a still image with voice-over versus using full motion video, for example, for particular purposes and learners. Hopefully, experts in cinema and curriculum design can merge techniques to construct effective and enjoyable learning experiences. Thus, a multimedia approach is important both to convey recorded content and to make effective interfaces. In the spirit of next generation "hyperknowledge" systems, the reader should link discussions of "multimedia" in the other papers in this collection to concepts of "knowledge" and "interface" occurring in the rest of this paper to remind us of this additional perspective.

A learning perspective implies exploring conceptual understanding as well as narrowing to practical application of knowledge, to include academic education, informal lifelong learning, and professional and industrial training. Besides providing help in mining information from librarylike repositories, an intelligent system must help in refining and creating knowledge—it should have many of the qualities of coach, tutor, and colleague to encourage the learner to question, conjecture, create, and experiment. Although an educational information system might help us find "the answer" more efficiently (still an important skill in information seeking), the emphasis must be on creating questions, proposing solutions, and contributing to understanding. Talking about "learners" instead of "users" should help make these information needs primary in our minds over information sources.

Adding intelligence to a multimedia educational information system could mean many things. We could build expert systems in many narrow domains, or expend effort to create an autonomous intellect equal or superior to our own. A more widespread interpretation of the meaning of artificial intelligence (AI) is to augment human intelligence via better development and authoring environments. Instead of regarding an intelligent system as a human replacement, we can consider the system as an adaptable tool, intelligent medium, or helpful assistant or partner [Lenat83; Stefik86; Fischer86]. Part of this means improving the organization of information—thus conveying "knowledge" embedded in a neighborhood of concepts rather than as isolated facts. Another aspect of intelligence is to "go beyond the information given" [Bruner73] to make inferences and interpretations. Intelligence in filtering or guiding means taking advantage of the relationships between concepts (for example, inheritance, ordering) to begin to guess what the learner understands (or misunderstands), to anticipate what he could do next or to prescribe a preferred path. To be intelligent, a system must grasp something about the purpose of the knowledge—what it can be used for—even if it does not deeply understand the information itself. To be intelligent, a system should be prepared to explain its actions (a general expert-system characteristic), but also be able to take advantage of explanations that the user offers. Learner and system should be able to converse at the level of examples, to develop analogies, and to propose generalizations.

Learner Intentions and Needs

In early attempts to supply electronic information, we have often modeled some of the artifacts such as books and libraries and little of the context that is such a vital part of the learning process. For example, a teacher, a written teacher's guide, or discussion group can add much to the understanding of a book. Card catalogs can be useful in a library—but wouldn't you rather have access to a knowledgeable librarian or research assistant?

What is the educational intent of searching? A helpful electronic assistant must know something beyond "the learner is collecting information" — something more like "the learner wants information to answer the following question or to help test the following hypothesis." A good tutor should have a grasp of a student's understanding and know why and when to offer knowledge that could be relevant or helpful. This is easier said than done. And yet it is a critical step in solving the problem that is of interest; simpler approaches typically give us answers to the wrong problems.

In some situations and for some learners, a quick, small dose of facts may be exactly what is desired. However, the information retrieved can suit many purposes depending on how it is presented. For example:

Tell me.	Give me the facts, no embellishments.
Inform me.	Facts plus optional background and other points of view.
Amuse me.	Find me interesting connections or perspectives.
Challenge me.	Make me find or create creative connections.
Guide me.	Let me browse but give over-the-shoulder advice.
Teach me.	More step-by-step guidance, fewer irrelevant links.

These needs, among others, are often not considered in information systems. Dervin [1983] succinctly contrasts "...two perspectives for looking at information transfer—one, an observer construction, sees information as a brick to be tossed from system to person [an empty bucket]; the other, a user construction, sees information as clay to be molded and shaped by the perceiver. The observer construction has been and still is the dominant perspective in the social science literature relating to information transfer and in the practice and design of information systems....In the context of such [absolute information] assumptions, an emphasis on information takes precedence over an emphasis on users. In the context of relativistic information assumptions, we would expect our information systems to look quite different. For one thing, the user would become the emphasis, the raison d'être. Instead of an emphasis on factualizing, there would be an emphasis on personalizing.... Instead of an emphasis on document transfer, there would be an emphasis on making sense."

Figure 1 illustrates the differences between these two perspectives on information systems and user activities. Current on-line systems are heavily biased toward the tidier, central view of information; hypothetical users of Memex would be encouraged to annotate, interpret, and personalize. As we begin to construct hyperknowledge systems giving access to large amounts of information for many different individuals, will we be able to maintain this balance between providing personalized access and efficient access? In addition to obtaining content, can we help learners understand how to request and interpret it? In addition to helping to better ask "What?" can we make more explicit the questions: "When is it appropriate?" "From whom did it come?" and "With what effect?"

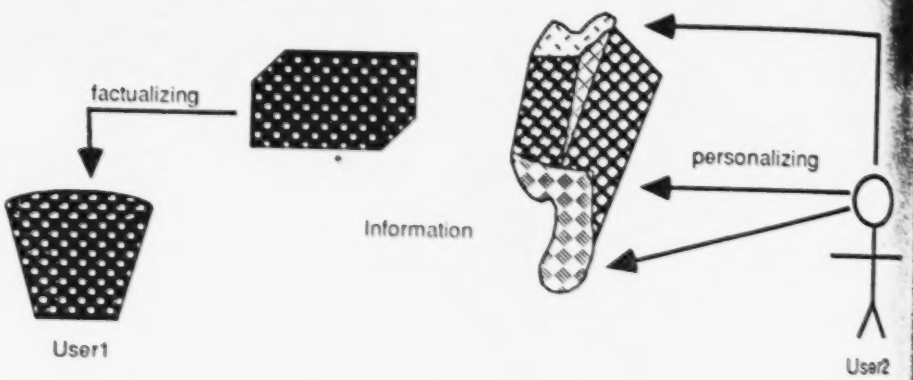

Figure 1.
Two perspectives on information systems: Absolute information "brick" delivered to user versus user creates meaning in information "blob"

Who is in the best position to articulate the needs and intentions of learners? Ideally, the learners themselves, even if inexperienced; however, what of the judgments of teachers, AI-based tutors, parents, the local community, or national or global committees? Choice of curriculum materials raises controversy when it touches on such issues as creationism vs. evolution, banning of certain novels for teenagers, sex education, and so on. Ironically, electronic information systems can provide broader access to information (assuming we can provide equal access and avoid the widening social and information gap between "have nets" and "have-nots") at the same time as they provide increased control for narrowing access. Censorship is really a form of filtering imposed by someone else that is difficult to be aware of or to undo. A key question is whether an electronic assistant should be somewhat subversive in making us aware that there are other points of view and the bias (perspective) of each, even if we have deliberately chosen to ignore them in our dedicated (or close-minded) attempts as entrepreneurs, scientists, religious seekers, or learners. At the same time, how can we avoid being deluged with multitudinous perspectives and opinions, some of which are more important or relevant to our needs than others? (See the later discussion of whether all links are created equal.) In many situations, the constraining guidance offered by editorial and tutorial assistants may best serve the learner otherwise overwhelmed by complete freedom of choice. Clearly, the dialog between system and learner to describe the right kind and right amount of information is never really complete, and must be periodically reexamined and refined.

In order to adapt to the learner's needs, the system will need to build a model of preferences, skills, and what the learner has previously looked at (how long ago, and how well understood). For example, what constitutes "news" is highly related to importance, uniqueness, and incremental value—the first several stories about the Chernobyl nuclear disaster are likely to be highly important regardless of my stated interests, thus overriding a narrow information filter. Assuming that additional stories are frequent and highly redundant in background information, a filter should then yield the significant differences or periodic summaries, not entire hourly stories. Many such "important" stories or subjects may compete for the learner's (or informed citizen's) attention. What is an important or unusual story or interesting interdisciplinary connection, though somewhat an editorial or author decision, is much more dependent on the learner's background.

Information vs. User-Oriented Organization

Some organizations of knowledge will certainly reflect conceptual relationships—multiple classifications of the world's knowledge (such as, Encyclopaedia Britannica's *Propaedia* or the CYC project [Lenat86]) to suit different cultural and learning styles. These correspond to the "observer constructs" or "absolute information" [Dervin83] discussed earlier. Conceptual relationships involving entities and events seem factual, for example, biological taxonomies (family of mammals), royal genealogies, sequences of events in world history (without necessarily implying cause-effect relationships), geographic (for example, France is in Europe, next to Belgium), and widely accepted scientific relationships (such as that between mass and energy).

Beyond the no-nonsense "tell me the facts" uses, lie pedagogical uses and organizations based more on "user constructs," that is, learner perceptions and needs. Pedagogical relationships between concepts admit to less of an absolute classification or agreement among experts; they correspond more to order and style of presentation, to personal interpretation, perception, and belief. In contrast to structuring the content itself, user-oriented organizations address questions of intent, function, appropriateness, source, and implications. For example:

These provide the following kinds of more detailed explanation of X.
Here is a diagram to summarize X.
Here is how Walter Cronkite might have described X.
These provide increasingly complicated examples of X.
Here are a series of experiments that can be used to prove (or disprove) X.
This argument supports proposition X under assumption Y.
This is a criticism/annotation of X.
Here are parameters to cause simulation X to fail in the following situation.
If X hadn't occurred, here's what might have happened.
How is person X similar to you?
This is analogous to X if we emphasize aspect Y.
This situation is similar to (different from) X in the following ways.

Dervin [1983] lists some possible related practices in a "user-based information system":

Frequent use of question-answer formats, as in the so-called "Dear Abby" format.
Frequent use of reporting values, motives, and frameworks for observing along with results of observing.
Frequent reporting of conflicting observations resulting from different frameworks, frequent use of case studies in presenting information.

In addition to organizing the information itself (the "answers"), a user-oriented approach should provide a structure for the possible approaches to the information (the "questions"). Simple lists (as above) are not satisfactory. Reverse links from information (as answers) to possible questions are not enough. We would like a question browser or problem browser that would relate types of questions and suggest appropriate strategies. How can kinds of questions be classified? Which typify the philosophic or pragmatic approaches of certain disciplines or individual practitioners? For example, what interesting questions would a biologist see in this situation, or how would Einstein (Newton, Galileo, Aristotle) have addressed this question? Creating a taxonomy of learner needs, styles, and problem-solving approaches may seem like an overwhelming task, but it is a necessary one if future systems are to aid the learner in more intelligently utilizing information, or to guide the learner.

An Example Question: Making Historical Comparisons

For example, if the learner has as a goal making comparisons by identifying similarities and differences, this suggests certain strategies for locating information, whether researching the best refrigerator or spreadsheet to buy or describing related work in a scientific survey article (for example, how is Memex like X?). In previous work [Weyer82], I gave students history problems such as "Compare the falls of the Roman and Han [Chinese] Empires," and provided electronic search and bookmark aids so that students could flip back and forth easily between the appropriate sections of an on-line history book. "The Huns attacked both empires" may not really be an adequate response to the question, even though it may be discussed for each empire chronologically and textually near the date of actual downfall. A more serious analysis would identify many earlier factors—social, economic, geographical—that may have led to weakness, and compare these across these empires in a static table format. An intelligent assistant might help us generalize the question to other empires or related problems. "Fall of an empire" is really a special case of "change in government" and "culture shock"—what factors differentiate it from civil war or scientific paradigm shifts? Beyond providing particular answer templates or tools to transform the question, an intelligent assistant might test deeper understanding by encouraging the student to try to influence the course of history, in games such as Chris Crawford's Eastern Front (an Atari 800 game to understand the failure of Hitler's invasion of Russia) or Balance of Power (a Macintosh game to explore political influence and intrigue in the modern nuclear world), or perhaps to create his or her own simulation of civilization in order to predict revolutions or declines.

Similarity-Based Searching

Searching in a user-oriented information system should be in terms of the learner's own concepts and words. In order to match phrasings not supplied explicitly by the user, natural language parsing, thesaurus tools, or graphical concept organizations must be available. Another possibility is that although the learner might be able to state a general description of the area of interest, in many situations a particular example is at hand, and the learner would like to find more "just like it, almost." For example, after reading about the fall of the Roman Empire, the learner might be interested in "similar" events. This may be easy if the author of the materials has provided explicit links to the declines of other empires—how much more powerful if these links could be declared in terms of the attributes the author considered important in suggesting the connection.

If adequate indexing has been done to identify significant concepts such as "Hun invasion," "political corruption," "mercenary army," and "immorality," then other events with some of these terms could be matched. However, manual indexing is time-consuming and often inconsistent, either for different trained indexers or even the same indexer over time. Automatic indexing is usually based on word frequency and can be improved by grouping words into index phrases; current systems do not collapse synonymous references to a single term, handle proper names properly, or suggest terms that do not occur in the document at all. Without indexing, it's easy to find a few relevant items using free text on most on-line systems but very difficult to be more exhaustive. However, a promising demonstration has been done on the massively parallel Connection Machine (64,000 processor machine from Thinking Machines), where most of the words in a document are used to match similar documents [Stanfill86].

There are several problems in finding similar items in current on-line systems. The problems of cumbersome syntax or differing query dialects between multiple information systems have been solved somewhat by search intermediaries or information gateways—programs with user interfaces to mask and map some of the complexity related to logging in, making queries (usually Boolean in syntax), and displaying results. Still, matching is too often based on exact occurrences of index terms or text words. Although manual use of word root truncation and thesauri of synonyms help to a limited extent, what is really required is a way of mapping these words into concepts located in semantic networks. Items that fail to match exactly may still be located within a small neighborhood of nearby concepts. It should be possible for the learner to define the radius of generality, specificity, and/or relatedness for a concept, as well as to set the level of importance in the overall combination of concepts used in the description (for example, corruption should probably include bribery and political favors, and correlate inversely with honesty; corruption is probably more important than lead poisoning in identifying similar situations).

Based on examples, it should be possible to define "interest neighborhoods" by locating each attribute of the example in the concept index and moving, narrowing, or expanding the focus of the concept. This is similar to the approach of the Rabbit system [Tou82]—query by reformulation, where the user refines a query by explaining what's wrong with an example. Alternatively, it might be possible to deal entirely with a collection of examples, with the user giving feedback like "good," "OK," or "bad" for each example as a whole without going into detail why—marking individual sentences, paragraphs, or articles as good examples of what the learner is interested in [Stanfill86]. If explanations of reasoning to the user are a hallmark of AI systems, then we should intuitively expect the reverse, that giving examples and explanations to the system could be used to demonstrate and debug the learner's intelligent behavior.

Beyond Browsing and Memex

Traditional systems often impose a sense of tunnel vision on the information—like accessing books in a library via card catalog and order form: you must guess at how it is described; when looking at the results of a search, it's often unclear what else of relevance that you missed. Hypertext systems now coming into use hold the promise of letting us look at the library shelves to perceive neighborhoods of information and many interesting connections within and between subject areas (for example, cross-references between encyclopedia articles, bibliographic references to books). Initially, the ability to see the relationships between many nuggets of knowledge is appealing. However, there are many problems to be overcome.

Without help, browsing everything can be as ineffective as searching blindly, thus creating a higher level of "information gridlock"—our electronic access is to much more information than was available in earlier generations, especially with technologies like CD-ROM storage and fiber optic networks. A large number of interesting associations between pieces of information will be available (and an overwhelming number of perhaps less interesting connections from the view that "everything is connected to everything else"). Browsing this universe directly without intelligent filtering to create smaller neighborhoods—or guidance to prioritize trail selection—will be frustrating.

Although there is much leverage to be gained by thinking of on-line information systems as "books" or "libraries," this metaphor can overshadow the dynamic nature of the information and its uses. We may think of the system more as pages to peruse rather than as a tutor to converse with or as a laboratory to conduct experiments in (for example, model-view-filter-guide metaphor [Weyer85]). To prevent "dusty-tome syndrome" (a disinterest or distaste for paperlike presentations), electronic "books" must become nonbooks in both function and name.

Knowledge is not static—to be user-oriented and malleable, it should be adaptable to the user's preferences. To illustrate this at the conference, I showed a videotape of a prototype electronic encyclopedia more fully described in [Weyer85]. Figure 2 shows text displayed in boldface to highlight active cross-references to other articles and provides conversions of measurements from English to metric units. Figure 3 shows a snapshot from an interactive simulation for the concept of harmonic motion; parameters can be adjusted via the sliders, and the spring animates.

Hyperknowledge systems should be able to generate new content via rules, look for inconsistencies, form hypotheses about pedagogy and understanding from the learner's actions, explain their own actions, and participate in the authoring process by classifying and critiquing individual annotations and revisions and refereeing debates by a community of learners (both human and artificial). Is there a real difference between author and reader, teacher and learner? As knowledge accumulates and ferments in the system, it should not all be treated equally. If the learner believes that knowledge that is heavily editorialized and filtered is better for his or her purposes, than, say, electronic mail discussions or political advertisements, then that value judgment should be an explicitly represented bias in the source selection.

All links are not created equal—they should depend on what you have seen before, on your current goals, your confidence in the sources, and on the author's guesses about what is most important for which of these contexts and purposes. Links need to be constantly annotated—why I find it useful, why I think someone else might be interested, why the system suggested it to me. The learner must communicate intentions so that a system can highlight and prioritize what's important, and hide connections and details considered to be unimportant (the learner could always reveal more and ask why some things were revealed and not others).

Figure 2

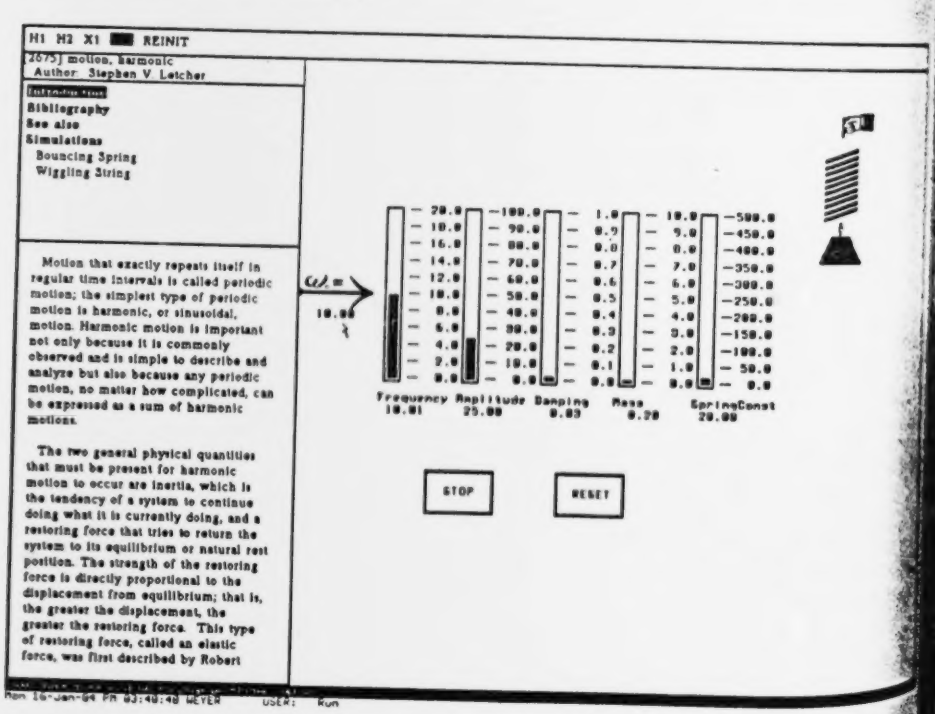
H1 H2 X1 REINIT
2675] motion, harmonic
Author: Stephen V. Letcher
Bibliography
See also
Simulations
Bouncing Spring
Wiggling String
Motion that exactly repeats itself in regular time intervals is called periodic motion; the simplest type of periodic motion is harmonic, or sinusoidal, motion. Harmonic motion is important not only because it is commonly observed and is simple to describe and analyze but also because any periodic motion, no matter how complicated, can be expressed as a sum of harmonic motions.
The two general physical quantities that must be present for harmonic motion to occur are inertia, which is the tendency of a system to continue doing what it is currently doing, and a restoring force that tries to return the system to its equilibrium or natural rest position. The strength of the restoring force is directly proportional to the displacement from equilibrium; that is, the greater the displacement, the greater the restoring force. This type of restoring force, called an elastic force, was first described by Robert
Frequency
Amplitude
Damping
Mass
SpringConst
STOP
RESET

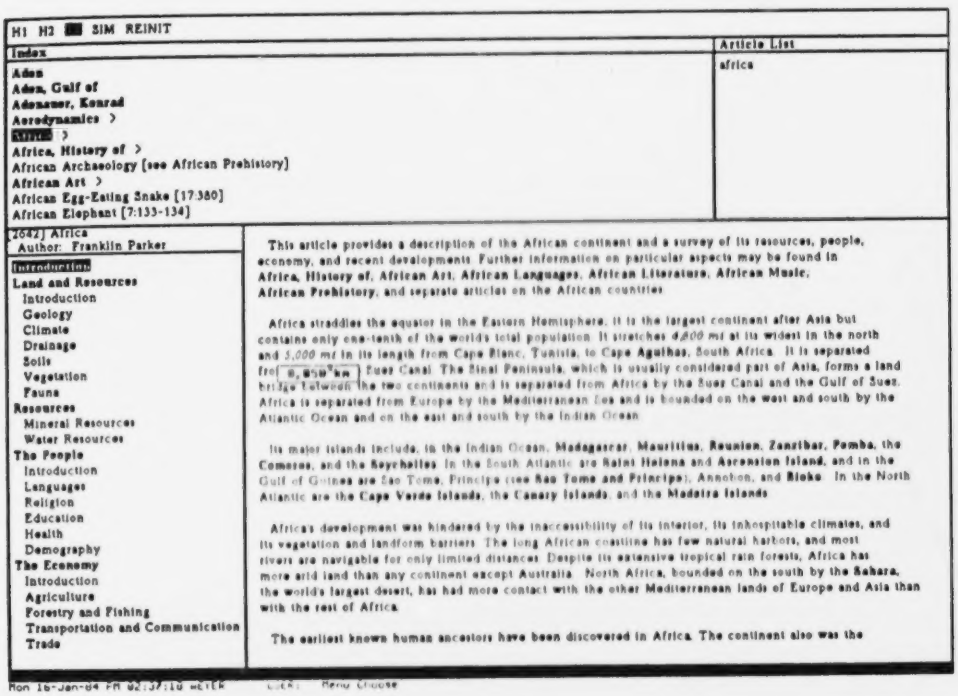
H1 H2 SIM REINIT
Index
Article List
africa
Adan
Aden, Gulf of
Adenauer, Konrad
Aerodynamics >
Africa, History of >
African Archaeology [see African Prehistory]
African Art >
African Egg-Eating Snake [17:380]
African Elephant [7:133-134]
[642] Africa
Author: Franklin Parker
Introduction
Land and Resources
Introduction
Geology
Climate
Drainage
Soils
Vegetation
Fauna
Resources
Mineral Resources
Water Resources
The People
Introduction
Languages
Religion
Education
Health
Demography
The Economy
Introduction
Agriculture
Forestry and Fishing
Transportation and Communication
Trade
This article provides a description of the African continent and a survey of its resources, people, economy, and recent developments. Further information on particular aspects may be found in Africa, History of; African Art; African Languages; African Literature; African Music; African Prehistory; and separate articles on the African countries.
Africa straddles the equator in the Eastern Hemisphere; it is the largest continent after Asia but contains only one-tenth of the world's total population. It stretches 4,600 mi at its widest in the north and 5,000 mi in its length from Cape Blanc, Tunisia, to Cape Agulhas, South Africa. It is separated fro Suez Canal. The Sinai Peninsula, which is usually considered part of Asia, forms a land bridge between the two continents and is separated from Africa by the Suez Canal and the Gulf of Suez. Africa is separated from Europe by the Mediterranean Sea and is bounded on the west and south by the Atlantic Ocean and on the east and south by the Indian Ocean.
Its major islands include, in the Indian Ocean, Madagascar, Mauritius, Reunion, Zanzibar, Pemba, the Comoros, and the Seychelles. In the South Atlantic are Saint Helena and Ascension Island, and in the Gulf of Guinea are São Tomé, Príncipe (see São Tomé and Príncipe), Annobón, and Bioko. In the North Atlantic are the Cape Verde Islands, the Canary Islands, and the Madeira Islands.
Africa's development was hindered by the inaccessibility of its interior, its inhospitable climates, and its vegetation and landform barriers. The long African coastline has few natural harbors, and most rivers are navigable for only limited distances. Despite its extensive tropical rain forests, Africa has more arid land than any continent except Australia. North Africa, bounded on the south by the Sahara, the world's largest desert, has had more contact with the other Mediterranean lands of Europe and Asia than with the rest of Africa.
The earliest known human ancestors have been discovered in Africa. The continent also was the

Mon 16-Jan-84 PM 02:37:10 WEIER

BEST COPY AVAILABLE

Guiding

Although work in search intermediaries and better browsing and hypertext techniques may improve searching and browsing in information systems, fundamental change is more likely to arise from trying to create intelligent assistants who know not only more about what we want to find but also why we want to find it. Natural language understanding has reached a similar barrier: new progress in understanding is not likely to come from further work in syntax and semantics, but from emerging theories of discourse. Discourse includes not only a history of the conversation (most natural language systems are impressive only on single utterances) but also the speaker's intentions and listener's expectations. Natural language itself could provide much of the required richness of dialog between learner and guide/assistant. Even if the user interface is not based on natural language, however, theories of discourse should influence its design.

Intelligence in an agent will not be synonymous with omniscience, even assuming that the learner can articulate what he or she is interested in finding and understanding. Although there can be a high correlation between an accurate description and relevant retrieved results, it is by the nature of communicating about information an imperfect match. There is a surprising discrepancy among human indexers about the appropriate terms to classify a document, though perhaps this can be improved with appropriate tools [Lenat86]. The phrasing of a query is not something that happens correctly the first time except in trivial cases—refinement and clarification are a natural part of the process. A resulting interest neighborhood will not contain all of the relevant items, and it will include irrelevant ones as well (that is, it will not exhibit perfect "recall" and "precision," in respective information jargon)—exploration, expanding, narrowing, highlighting, and rejecting should be expected. In trying to make our assistant emulate expert human behavior, we must realize that this seeming imperfection in communication is normal and, indeed, essential in modeling learning.

Conclusion

The powerful vision typified by Memex can be made even more compelling by transforming information into knowledge, by elevating the needs and importance of the learner, and by evolving a flexible tool in the direction of becoming an adaptive partner. Multimedia techniques are important not only to convey detail, motion, and emotion of knowledge content but also to create new organization and navigation interfaces to access this knowledge. Even if we regard our "hyperknowledge assistant" as only an information system, the preceding discussion should lead us at least to new ways of thinking about searching and browsing. If we consider it additionally as an exploratory and guided learning environment, learner needs and questions must become the focus.

Next steps? Many are the same as in [Weyer85]; we need more experience in highly interconnected hypermedia systems, intelligent tutoring systems, knowledge organizations, programming and authoring environments, conceptual and example-based description, natural language discourse and understanding systems, and models of user intentions, understanding, and behavior. Integrating these pieces will be a major conceptual, as well as software engineering, task. It is my hope that the vision I have attempted to present here can be elaborated and made practical before another 40 years have elapsed.

References

Bruner, J. *Beyond the Information Given.* New York: W.W. Norton, 1973.

Bush, V. As We May Think. *Atlantic Monthly*, 176: 101-108; 1945 July. Also reprinted in Lambert S. and Ropiequet, S. eds. *New Papyrus*, Microsoft Press, 1986, pp. 3-20.

Dervin, B. Information as a User Construct: The relevance of perceived information needs to synthesis and interpretation. In S. A. Ward and L. J. Reed (Eds.) *Knowledge Structure and Use: Implications for Synthesis and Interpretation.* Temple University Press, Philadelphia, PA, 1983, pp. 153-183.

Fischer, G. Enhancing Incremental Learning Processes with Knowledge-Based Systems. In H. Mandl and A. Lesgold (Eds.) Learning Issues for Intelligent Tutoring Systems, Springer-Verlag, New York; 1986 (in press).

Lenat, D., Prakash, M., Shepherd, M. CYC: Using Common-Sense Knowledge to Overcome Brittleness and Knowledge-Acquisition Bottlenecks. *AI Magazine.* Vol. 6. No. 4, pp. 65–85. Winter 1986.

Lenat, D., Borning, A., McDonald, D., Taylor, C., and Weyer S. Knoesphere: Building Expert Systems with Encyclopedic Knowledge. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence* (Karslruhe, West Germany). 1983, pp. 167–169.

Stanfill, C. and Kahle, B. Parallel Free Text Search on the Connection Machine System. Thinking Machines Corp. Technical Report No. 86-5, March 1986.

Stefik, M. The Next Knowledge Medium. *AI Magazine.* Vol. 7, No. 1, pp. 34–46. Spring 1986.

Tou, F., Williams, M., Fikes, R., Henderson, D.A., and Malone, T. RABBIT: An Intelligent Database Assistant. In *Proceedings of the American Association for Artificial Intelligence* (Pittsburgh). August 16–20, 1982.

Weyer, S. A. The Design of a Dynamic Book for Information Search. *International Journal of Man Machine Studies.* Vol. 17, No. 1, July 1982, pp. 87–107.

Weyer, S. A. and Borning, A. H. A Prototype Electronic Encyclopedia. *ACM Transactions on Office Information Systems*, Vol. 3., No. 1, Jan. 1985, pp. 63–88.

Stephen Weyer

Stephen Weyer has worked for 16 years in enhancing learning environments on computers, improving access to knowledge, and providing advanced programming tools. At Stanford, he developed and tested Logo environments involving animation, sound, and model trains. Over the next 9 years at Xerox P.A.R.C., he taught and extended Smalltalk, and developed information retrieval tools, browsing interfaces, and applications for the office, library, and school. For his dissertation work at Stanford, he designed an on line world history book and evaluated its use by students to answer questions. In contrast to this emphasis on structure and search techniques, his work on a prototype electronic encyclopedia at Atari explored issues related to dynamic information: active objects, videodisc images, and simulations. During two years at Hewlett-Packard, he managed an AI programming environment project (a frame-based knowledge browser). At Apple, he is balancing and combining these themes of exploratory and guided learning environments, information browsing and creation, and artificial intelligence for representing and presenting knowledge.

21